MySQL 8.0.33 릴리즈 노트

MySQL 8.0.33 변동 사항 (2023-04-18 GA)

감사 로그 참고 사항

- 이전에 MySQL Enterprise Audit는 필터 및 사용자 계정 데이터의 영구 저장을 위해 mysql 시스템 데이터베이스의 테이블을 사용했습니다. 유연성을 높이기 위해, 새로운 audit_log_database 서버 시스템 변수를 사용하면 서버를 시작할 때 글로벌 스키마 네임스페이스의 다른 데이터베이스를 지정할 수 있습니다. 테이블 스토리지의 기본값은 mysql 시스템 데이터베이스입니다. (WL#15500)

컴파일 참고 사항

- **마이크로소프트 윈도우**: Visual Studio 2017 이상의 MSVC 코드 분석 지원이 추가되었습니다. 이렇게 하면 현재 디렉터리와 해당 하위 디렉터리에서 이 분석을 활성/비활성하도록 설정하는 MSVC_CPPCHECK(기본값은 OFF) CMake 옵션이 추가됩니다. (버그 #34828882)
- MySQL이 GNU 컴파일러 또는 clang으로 빌드되며 curl 버전이 7.86보다 큰 경우, curl 대체 예정 경고가 -Wno-error로 다운그레이드되었습니다. (버그 #35111625)
- macOS에서 필요에 따라 공유 시스템 라이브러리와 링크하기 위해 curl 인터페이스를 연결할 때 -framework CoreFoundation 및 -framework SystemConfiguration을 추가했습니다. (버그 #35104962)
- MY_INCLUDE_SYSTEM_DIRECTORIES 매크로를 라이브러리 인터페이스로 대체했습니다. (버그 #35018072, 버그 #35028089, 버그 #35072295)
- 대체 링커를 지원하도록 CMake 코드를 개선했습니다. (버그 #34963568)
- 더이상 사용되지 않는 Docs/mysql.info 파일을 빌드 시스템에서 제거했습니다. (버그 #34960126)
- 최상위 .clang-tidy 파일과 관련된 .clang.tidy 파일을 strings/ 디렉토리 및 mysys/ 디렉토리에 추가했습니다. 또한 compdb 지원을 활성화하여 헤더 파일에서 clang-tidy 사용을 활성화했습니다. (버그 #34917075)
- uca-dump나 uctypedump와 같은 기능들에 유지 보수되지 않거나 사용되지 않은 C++ 소스 파일을 일부 삭제했습니다. (버그 #34898978)
- 커맨드 라인에서 컴파일 할 때 GCC 및 Clang의 컬러 컴파일러 출력을 활성화하는 CMake 빌드 옵션이 추가되었습니다. 활성화하려면 -DFORCE_COLORED_OUTPUT=1을 CMake에 전달합니다. (버그 #34897192)
- 윈도우의 경우, 써드 라이브러리에서 관련된 .dll 파일이 발견되면 해당 .pdb 파일도 설치합니다. (버그 #34863555)
- 엔터프라이즈 리눅스 8 빌드 및 엔터프라이즈 리눅스 9 빌드는 이제 GCC 11 대신 GCC 12를 사용합니다. (버그 #34829151)
- -static-libgcc -static-libstdc++를 사용하여 빌드하면, 필요에 따라, 정적 라이브러리를 사용하여 번들된 프로토콜 버퍼도 같이 빌드합니다. (버그 #110216, 버그 #35123848)

컴포넌트 참고 사항

- 이제 INSTALL COMPONENT는 하나 이상의 컴포넌트를 설치하는 동안 컴포넌트 시스템 변수의 값을 설정하는 SET 절이 추가됩니다. 이 새로운 절은 변수 값을 할당하는 다른 방법과 관련된 불편과 제약점을 줄여줍니다. 사용법에 대한 자세한 내용은, INSTALL COMPONENT 문을 참조하세요. (WL #10916)

사용 중단 및 제거된 기능 참고 사항

- 사용자 정의 콜레이션(문자 집합에 콜레이션 추가하기 참고)은 더 이상 사용되지 않습니다. 다음 중 하나를 사용하면 경고가 로그에 기록됩니다:
  - SQL 문에서 COLLATE 다음에 사용자 정의 콜레이션명을 넣는 경우
  - collation_server, collation_database 또는 collation_connection의 값으로 사용자 정의 콜레이션명을 사용하는 경우

  사용자 정의 콜레이션에 대한 지원은 향후 MySQL 버전에서 제거될 예정입니다. (WL #14277)

MySQL Enterprise 참고 사항

- MySQL Enterprise Edition은 이제 MySQL 8.0.13에서 도입된 플러그인 라이브러리를 기반으로 하지 않고 컴포넌트 기반 데이

터 마스킹 및 익명화 기능을 제공합니다. 컴포넌트 구현은 사전을 관리하기 위한 전용 권한을 제공하고 다음을 포함하는 특정 유형의 목록을 확장합니다.

- Canada Social Insurance Number(캐나다 사회 보험 번호)
- United Kingdom National Insurance Number(영국 국민 보험 번호)
- International Bank Account Number(국제 은행 계좌 번호)
- Universally Unique Identifier (UUID 범용 고유 식별자)

향상된 테이블 기반 사전 레지스트리는 플러그인에서 사용하는 파일 기반 사전을 대체합니다. 컴포넌트와 플러그인 구현 간의 차이점에 대한 요약은 데이터 마스킹 컴포넌트와 데이터 마스킹 플러그인 비교를 참조하십시오. 기존 플러그인 사용자는 새로운 MySQL Enterprise 데이터 마스킹 및 익명화 컴포넌트를 설치하기 전에 서버측 플러그인을 제거하고 로드 가능한 기능을 제거해야 합니다. (버그 #33851601, WL #12641)

Performance Schema 참고 사항

- 이번 릴리즈에는 Performance Schema Server Telemetry Traces 서비스가 추가되었습니다. SQL문의 수명과 관련된 알림을 검색할 수 있는 방법을 플러그 인과 컴포넌트에 제공하는 인터페이스입니다.
  아래 항목들이 추가되었습니다:

  - 상태 변수 Telemetry_traces_supported. 서버 원격 측정 추적이 지원되는지 여부. (Boolean 값)
  - TELEMETRY_ACTIVE 열이 threads 테이블에 추가되었습니다. 스레드에 활성 원격 측정 세션이 연결되어 있는지 여부를 나타냅니다.

  (WL #15059)

추가 또는 변경된 기능

- **중요한 변경 사항**: OpenSSL 라이브러리가 번들로 제공되는 플랫폼의 경우, 링크된 MySQL 서버용 OpenSSL 라이브러리가 버전 1.1.1t로 업데이트되었습니다. OpenSSL 버전 1.1.1t에서 수정된 문제는 https://www.openssl.org/news/cl111.txt에 설명되어 있습니다. (버그 #35092429)
- **복제** : MySQL 제품에서 사용되는 오래된 용어들을 변경하기위해 진행중인 작업의 일환으로 "master", "slave", "MTS"라는 용어를, MySQL 복제 관련된 오류 메시지에서, source”, “replica”, “MTA”로 대체하였습니다. 여기에는 messages_to_clients.txt 및 messages_to_error_log.txt에 나열된 복제와 관련된 모든 오류 메시지가 포함됩니다; 현재 작업에서는 (에러 메세지가 아닌 / 복제 관련이 아닌) 다른 컨텍스트에서 사용되는 메세지에 대해서는 변경을 수행하지 않았습니다.
  자세한 내용은 MySQL 8.0 오류 메시지 레퍼런스를 참조하세요. (버그 #108422, 버그 #34594819, WL #14191)
- **복제**: **mysqlbinlog** --start-position은 이제 최대 18446744073709551615까지의 값을 허용합니다. 그러나 --read-from-remote-server 또는 --read-from-remote-source 옵션이 같이 사용되는 경우 최대값은 4294967295입니다. (버그 #77818, 버그 #21498994)
- 시스템의 curl 라이브러리에 링크하는 대신, curl을 포함하고 있는 바이너리 패키지의 경우 curl 버전 7.88.1을 사용하도록 업그레이드되었습니다. (버그 #34828111)
- 디폴트 값으로 DEFAULT(col_name)이란 형태로 생성하여 기본값으로 컬럼명을 지정 하는 것은 이제 허용되지 않으며 오류 메시지가 출력됩니다. (버그 #34463652, 버그 #34369580)
- 이제새로운 TELEMETRY_LOG_ADMIN 권한으로 원격 분석 로그 구성을 활성화할 수 있습니다. 이 권한은 MySQL HeatWave on AWS를 통해 독점적으로 배포되는 telemetry_log 플러그인에 의해 정의됩니다. (Bug #111395, Bug #35494180)
- 이제 성공적으로 실행된 명령문에 대하여 Statement ID를 클라이언트에게 반환할 수 있습니다. 세션별로 이 기능을 사용하려면 session_track_system_variables의 값으로 statement_id를 주거나, session_track_system_variables의 값으로 특별 값인 * (아스테리스크)을 지정합니다. (WL #15418)

버그 수정

- **NDB Cluster:** Occasional temporary errors which could occur when opening a table from the `NDB` dictionary while repeatedly performing concurrent schema operations were not retried. (Bug #34843889)
- **NDB Cluster:** During iteration, ordered index scans retain a cursor position within each concurrently scanned ordered index fragment. Ordered index fragments are modified and balanced as a result of committing DML transactions, which can require scan cursors to be moved within the tree. When running with query threads configured (`AutomaticThreadConfig` set to 1), multiple threads can access the same index fragment tree structure, and the scans of multiple threads can have their cursors present in the same structure.
  The current issue arose due to an assumption in the logic for moving scan cursors when committing DML operations that all scan cursors belonged to the LDM thread owning the index fragment, which did not allow for the possibility that such fragments might belong to query threads. (Bug #33379702)
  References: See also: Bug #32257063.
- **InnoDB:** Dead code removal. (Bug #35036850, Bug #109873)

- **InnoDB:** Error messages related to `innodb_doublewrite` moved to the error log. (Bug #34883045, Bug #109330)
- **InnoDB:** Prevent online DDL operations from accessing out-of-bounds memory. (Bug #34750489, Bug #108925)
- **InnoDB:** `ALTER TABLE ... AUTO_INCREMENT` could be set to less than MAX + 1 and not forced to MAX + 1. (Bug #33419246, Bug #105092)
- **InnoDB:** `Innodb_data_pending_fsyncs` could show extremely high inaccurate values because of a variable overflow. (Bug #30133150)
- **Partitioning:** Some `IN()` queries on partitioned tables were not always handled correctly. (Bug #34801284)
  References: This issue is a regression of: Bug #32311183.
- **Partitioning:** Queries using the `INDEX_MERGE` optimizer hint was not handled correctly in all cases. (Bug #34797257)
- **Replication:** XA transactions whose XIDs contained null bytes could not be recovered. (Bug #34918985)
- **Replication:** When `binlog_order_commits` was set equal to `1`, for any two transactions and for any sub-step of the commit phase, the transaction that was written to the binary log first did not always execute the sub-step first, as expected. (Bug #34703698)
- **Replication:** Some binary log events were not always handled correctly. (Bug #34617506)
- **Replication:** The binary log recovery process did not report all possible error states. (Bug #33658850)
- **Replication:** Following `CHANGE REPLICATION SOURCE TO SOURCE_CONNECTION_AUTO_FAILOVER=1`, failover generated a number of misleading warnings in the log that implied there were problems when in fact conditions were those expected for such a failover. These log messages have been updated accordingly. (Bug #32135376)
- **Replication:** When a transaction failed, as a side effect, extraneous error messages relating the replication data repositories were written to the log. Now in such cases, we suppress such error messages, which are not directly related to the issue of the failed transaction or its cause. (Bug #19820134)
- **Replication:** Setting `binlog_order_commits` to `OFF` could lead to a missed GTID in the next binary log file's `Previous_gtids` event.
  Our thanks to Yewei Xu and the Tencent team for the contribution. (Bug #109485, Bug #34930969)
- **Replication:** Corrected the SQL statements suggested in the error message text for `ER_RPL_REPLICA_ERROR_RUNNING_QUERY`.
  Our thanks to Dan McCombs for the contribution. (Bug #109154, Bug #34822612)
- **Replication:** A hash scan builds a hash of changes, scans the target table or index, and applies any matching change for the current entry. In the build phase, it uses only the before image, and skips any after image. Problems arose in some cases because generated columns were computed for the (skipped) after image, leading to replication errors. This is fixed by not computing generated columns any longer for seek-only calls such as hash scans.
  Our thanks to dc huang for the contribution. (Bug #107366, Bug #34198907)
- **Replication:** In certain rare cases, it was possible to set `gtid_mode=OFF` for one session while another session, after `WAIT_FOR_EXECUTED_GTID_SET()` was issued by a user in this second session, was still waiting for the next GTID set from the first session. This could result in the second session waiting indefinitely for the function to return. (Bug #99921, Bug #31505993)
- **Group Replication:** Accessing the Performance Schema `replication_group_communication_information` and `replication_group_member_stats` tables in parallel sometimes caused subsequent group replication operations to hang. (Bug #34870181)
- **Group Replication:** In certain cases, the group replication secondary node unexpectedly shut down while purging the relay log. (Bug #34397106)
- **Group Replication:** When shutting down the Group Replication plugin, the order in which the associated events were reported the error log sometimes led to confusion. To remove any doubts, we now make sure that `Plugin group_replication reported: 'Plugin 'group_replication' has been stopped.` is in fact the last log message relating to the shutdown, written only when all other events associated with shutting down the plugin have been logged. (Bug #109345, Bug #34887491)
- **Microsoft Windows:** The `authentication_fido_client` plugin stopped responding during the authentication process if it was unable to find a FIDO device on the Windows client host. (Bug #34918044)
- In certain cases, `CONVERT( utf8mb3_column USING UTF16)` was rejected with the error Cannot convert string '\x--...' from binary to utf16. (Bug #35129361)
- When joining two tables on a string column, and the column from one of the tables has an additional predicate comparing it with a temporal literal, constant propagation in some cases incorrectly caused the join condition to be modified such that it used temporal rather than string semantics when comparing the strings. This caused incorrect results to be returned from the join. (Bug #35115909)
- Error messages returned after calling the `mysql_reset_connection()` C API function in a prepared statement did not identify the function name properly. (Bug #35107280)
- Fixed a regression in a previous fix for an issue with windowing functions.
  Our thanks to Dmitry Lenev for the contribution. (Bug #35061924)
  References: This issue is a regression of: Bug #34572136.
- When replacing subqueries in transforms, the internal flag showing whether a given query block contains any subqueries (`PROP_SUBQUERY`) was not updated afterwards. (Bug #35060385)
- A client setting the character set to an impermissible client character set (`ucs2`, `utf16`, `utf16le`, or `utf32`) could cause unexpected behavior when the client used an authentication plugin. (Bug #35054579)
- `EXPLAIN ANALYZE` displayed 0 when the average number of rows was less than 1. To fix this, we now format numbers in the output of `EXPLAIN ANALYZE` and `EXPLAIN FORMAT=TREE` such that numbers in the range 0.001-999999.5 are printed as decimal numbers, and numbers outside this range are printed using engineering notation (for example: `1.23e+9`, `934e-6`). In addition, trailing zeroes are no longer printed, and numbers less than `1e-12` are printed as `0`.
  This helps ensure consistent precision regardless of the number's value and improve readability, while producing minimal rounding errors. (Bug #34990948)
- The `NTILE()` function did not work correctly in all cases. (Bug #34986665)
- Some joins on views did not perform correctly. (Bug #34985359)
- Transforming a correlated scalar subquery to a derived table led to a wrong result for `InnoDB` tables when the subquery included duplicate predicates. An example of a query which could be affected by this issue is shown here:

```
SELECT * FROM t1
WHERE (
```

```
SELECT t2.a FROM t2
WHERE t2.a = t1.a AND t2.a = t1.a
 ) > 0;
```

(Bug #34973220)

- Fixed an assert in `sql/item_strfunc.cc` that could potentially lead to issues with the `SPACE()` function. (Bug #34962821)
- Using `ROW_COUNT()` as the *length* argument to `LPAD()` or `RPAD()` did not perform as expected. (Bug #34961236)
- A query with a window function having an expression with a `CASE` function in its `ORDER BY` clause could lead to a server exit. (Bug #34933045)
- The fix for a previous issue introduced an assertion in debug builds when optimizing a `HAVING` clause. (Bug #34923792)
  References: This issue is a regression of: Bug #33725415.
- When using `mysqld_multi`, the system that obscures "--password" usage as "--password=*****" would also match "--password-history" and "--password-require-current" definitions as "--password", but now explicitly checks for "--password=" instead. (Bug #34918740)
- In some cases, calling the `mysql_bind_param()` C API function could cause the server to become unresponsive. (Bug #34869076)
- The `authentication_oci_client` plugin was unable to open a valid configuration file if any of its entries contained an equals sign character separated by spaces (for example, `key_file = /home/user/.oci/oci_api_key.pem`). Now, both '*key*=*value*' and '*key* = *value*' entry formats are supported. (Bug #34864078)
- Incorrect results were returned when the result of an `INTERSECT` or `EXCEPT` operation was joined with another table. This issue affected these operations in such cases when used with either `DISTINCT` or `ALL`. (Bug #34843764)
- When preparing a view query, the operation used the system character set (instead of the character set stored in data dictionary) and then reported an invalid character-string error. (Bug #34800905)
- Prepared statements that operate on derived tables, including views, could stop unexpectedly due to problems with the code for reopening tables after an error. (Bug #34798403)
- Removed an assertion raised in certain cases by the `RANDOM_BYTES()` function in debug builds. (Bug #34781507)
- There was an issue in how persisted variables were set on startup, causing certain variables not to get properly set to their persisted value. (Bug #34751419)
- The `MAKETIME()` function did not perform correctly in all cases. (Bug #34745241)
- Some functions with multiple arguments did not produce the expected results. (Bug #34741801)
- A table reference in an `ORDER BY` outside the parenthesized query block in which the table was used, and which query block had no `LIMIT` or `ORDER BY` of its own, raised an error. (Bug #34704011)
  References: This issue is a regression of: Bug #103954, Bug #32980249.
- A left join with an impossible condition as part of an `ON` clause was not optimized as in MySQL 5.7, so that in MySQL 8.0, the query executed more quickly without the impossible condition than with it. An example of such a query, impossible condition included, is `SELECT * FROM t1 JOIN t2 ON t1.c1=t2.c1 AND 1=2`. (Bug #34668756)
- When a user defined function was part of a derived table that was merged into the outer query block, or was part of a subquery converted to a semi-join, knowledge of whether this UDF was deterministic (or not) was lost during processing. (Bug #34666531)
- With JSON logging enabled and an event subclass specified in the audit log filter definition, an empty item `"":{}` was appended to the end of the logged event. (Bug #34659904)
- Some subqueries did not execute properly. (Bug #34616553)
- After the `asymmetric_encrypt()` component function in a `SELECT` query encountered a NULL field to decrypt, it could return NULL values for other non-NULL encrypted fields. (Bug #34598912)
- The server did not always shut down cleanly after uninstalling the audit log plugin. (Bug #34594035)
- Certain antijoins were not handled correctly by the server. (Bug #34370673)
  References: This issue is a regression of: Bug #30573446.
- When the MySQL 5.7 Optimizer has 2 choices for an index to filter rows, one primary and one secondary, it picks a range scan on the secondary index because the range scan uses more key parts. MySQL 8.0 did not use this logic, instead choosing the primary index to filter rows with `WHERE` clause filtering. Primary key use is not suitable in such cases due to the presence of `LIMIT`, and due to the nature of data distribution. The secondary index was not considered while resolving order by due to constant elimination. This resulted in much different query plans in MySQL 5.7 and MySQL 8.0 for the same query.
  We solve this issue in MySQL 8.0 by skipping the constant key parts of the index during order-by evaluation only if the query is constant-optimized, which can be done at this time, but not during `LIMIT` analysis. (Bug #34291261)
- The MySQL data dictionary caches failed lookups of `se_private_id` values (IDs which are not found), which speeds up execution of code specific to `InnoDB`, relying on the fact that `InnoDB` does not reuse these IDs. This assumption does not necessarily hold for other storage engines, most notably `NDB`, where this problem was resolved previously by not using this cache.
  We extend the previous fix made for `NDB` so that the cache lookup is now employed only when the table uses the `InnoDB` storage engine. (Bug #34145006)
  References: See also: Bug #33824058.
- Unexpected results were seen in some queries using `DENSE_RANK()`, possibly with the addition of `WITH ROLLUP`. (Bug #34099408)
- Fixed an assert raised in `sql/sql_tmp_table.cc` following work done previously to reimplement `ROLLUP` processing. (Bug #33830659)
  References: This issue is a regression of: Bug #30969045.
- Some CTEs that did not use any tables were not always handled correctly. (Bug #33725542)
  References: This issue is a regression of: Bug #27062031.
- Accessing rows from a window frame of a window function call present only in the query's `ORDER BY` list raised an error. (Bug #33069747)
- `PERCENT_RANK()` used with `ORDER BY` *column* did not return the correct result. (Bug #33064174)
  References: This issue is a regression of: Bug #27484133.
- The `--exclude-tables` and `--include-tables` **mysqlpump** options did not handle views. (Bug #21303549)

- Changed the MySQL systemd service unit configuration from `After=network-online.target` to `Wants=network-online.target` to ensure that all configured network devices are available and have an IP address assigned before the service is started. (Bug #109996, Bug #35068274)
- `AVG(...) OVER (ROWS BETWEEN 1 FOLLOWING AND UNBOUNDED FOLLOWING)` did not return the correct result. (Bug #109725, Bug #35013880)
  References: This issue is a regression of: Bug #108008, Bug #34431996.
- A query of the form `SELECT 1 FROM t1 WHERE NOT EXISTS (VALUES ROW(1), ROW(2))` caused an assert in debug builds when the `subquery_to_derived` optimizer switch was enabled. (Bug #109723, Bug #35014318)
  References: See also: Bug #108910, Bug #34746261.
- **mysqlimport** did not escape reserved word table names when used with the `--delete` option. (Bug #109711, Bug #34999015)
- When cloning a condition to push down to a derived table, characters in strings representing conditions were converted to `utfmb4` correctly only for values less than 128 (the ASCII subset), and code points outside the ASCII subset were converted to invalid characters, causing the resulting character strings to become invalid. For derived tables without `UNION`, this led to problems when a column name from the derived table used characters outside the ASCII subset, and was used in the `WHERE` condition. For derived tables with `UNION`, it created problems when a character outside the ASCII subset was present in a `WHERE` condition.
  We fix these issues by initializing the string used for representing the condition in such cases to the connection character set. (Bug #109699, Bug #34996488)
- Using `--single-transaction` with **mysqldump** version 8.0.32 required either the `RELOAD` or `FLUSH_TABLES` privilege. This requirement now applies only when both `gtid_mode=ON` (default `OFF`) and with `--set-gtid-purged = ON|AUTO` (default `AUTO`). (Bug #109685, Bug #109701, Bug #34993824, Bug #34998910, Bug #35020512)
  References: This issue is a regression of: Bug #105761, Bug #33630199.
- Fixed a number of issues present in the internal documentation for the scramble generator algorithm in `sha256_scramble_generator.cc` and `sha2_password_common.cc`.
  Our thanks to Niklas Keller for the contribution. (Bug #109576, Bug #34967141)
- `CREATE USER IF NOT EXISTS` added a password history entry even when the user already existed and the password was not updated. This caused a subsequent `ALTER USER` statement to be rejected. (Bug #109415, Bug #34906592)
- Many joins using `eq_ref` access did not perform as well as in previous versions. This issue was first reported in MySQL 8.0.29. (Bug #109361, Bug #34891365)
- A hash outer join sometimes incorrectly matched `NULL` with a decimal zero or an empty string that used a non-padding collation, leading to erroneous results. (Bug #109211, Bug #34837464)
  References: This issue is a regression of: Bug #33794977.
- An object used internally by `ALTER INSTANCE RELOAD TLS` was not freed until the number of readers reached 0, under the assumption is that the number of readers should reach 0 fairly frequently. The read lock held during an SSL handshake is generally an expensive operation, with network calls, so when roundtrips between the client and the server took excessively long, the lock was held for a relatively long amount of time. This meant that, when changing the value of this object and there were a sufficient number of incoming SSL connections being made, the number of readers might not reach 0 in a reasonable length of time, leaving the thread holding the lock using 100% of the CPU until the lock was released.
  We fix this by adding a wait after setting the pointer to this object to a new value, but before releasing the old object.
  Our thanks to Sinisa Milivojevic for the contribution. (Bug #107567, Bug #34284186)
- If **mysqldump** or **mysqlpump** could not convert a field's default value to UTF-8 (for instance, if the field was of type `BINARY` and the default value did not coincide with valid UTF-8), the operation produced results that were not valid to import. Further, using the `--hex-blob` option did not resolve the issue. We now convert the default value to the system character set. If this fails, the server sends the value as hexadecimal instead to make it more human-readable. (Bug #104840, Bug #33322551)
- A connection using the C API (`libmysqlclient`) client library could fail with the `FUTURE` crypto policy. (Bug #104778, Bug #33316709)
- While cloning a temporary table for a common table expression which used shared materialization, the cloned temp table was not marked as using hash deduplication, leading to wrong results. We now set the hash field for the cloned temporary table correctly, and update the hidden field count to take this into account. (Bug #102251, Bug #32383712)
  References: See also: Bug #103052, Bug #32659629.
- `CREATE EVENT` and `ALTER EVENT` assumed that all values passed to them (other than in a `DO` clause) resolved as scalars without actually checking the values. This led to assertions when any such values actually rows. We now perform an explicit check for the number of columns when resolving such items, and report an error when one produces a row and not a scalar value. (Bug #57552, Bug #11764690)
- A view reference whose underlying field is a constant is not marked as constant when the reference is part of an inner table of an outer join. It was found that, when pushing a condition down to a derived table, the reference was stripped off and only the underlying field was cloned, which made it a constant, and led to wrong results.
  To fix this problem, we ensure that we do not push such a condition down to the derived table by adding a check to see first whether the table used by the condition matches the derived table or is a constant expression; only when it is one or the other of these do we actually push the condition down. (Bug #34661, Bug #11747971)

Revision #4
★Created 8 September 2023 07:03:25 by 신민항
Updated 20 September 2023 04:47:24 by 신민항